Model-Based Controls for Integrated Shading and UFAD

B. Coffey

PhD Candidate, Architecture (Building Science), UC Berkeley

E. S. Lee

Building Technologies Program, Environmental Energy Technologies Division, Lawrence Berkeley National Laboratory, Mailstop 90-3111, 1 Cyclotron Road, Berkeley, CA 94720 USA

October 6, 2011 (with minor edits by B. Coffey, January 9, 2014)

Abstract

Methods and initial results are described for model-based controls with offline optimization for integrated shading and UFAD control for an office building in New York. Two cases are studied through lookup-table calculation and annual simulation of the resulting controllers: one case with interior blinds, the other case with exterior blinds. The interior blind case was found to reduce HVAC energy by 5% over a simple baseline control, and the exterior blind case produced a 5.6% HVAC energy savings over the baseline. Further investigations and case studies are planned.

Keywords: Model predictive control (MPC), offline optimization, cloud computing, integrated controls, shading and HVAC controls.

1. Background

A new and highly efficient office building in Manhattan was used for this study. The building has automated shading, automated dimmable lighting and an underfloor air distribution (UFAD) system. The automated shades and HVAC system are currently being controlled independently. The question at hand is whether there might be any energy savings available through integrated control of the two systems. The study also hopes to address this question more broadly than for just this particular building by considering different climates and variants on the system configurations.

2. Problem definition

Figure 1 shows the two control setpoints considered in this study: solar shading position and supply air temperature setpoint. The goal is to determine the optimum values for these setpoints (as a vector, noting their interdependence) that minimize combined HVAC and lighting energy consumption, for any given weather and loads conditions.

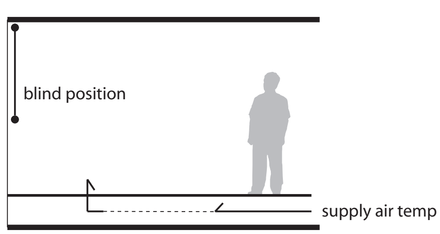

Fig. 1. Blinds and HVAC control variables of interest

The optimization must consider various complexities, including:

- the general tradeoffs between thermal loads and lighting energy for different blind positions
- the effects of solar gains to the floor on thermal decay in the UFAD system, which effects supply air flow rates and feasible supply air temperature ranges, and thus cooling COP and the potential for economizer operation
- capturing the solar gains at the blind creates a plume that drives thermal stratification in the zone, which affects cooling COP.

Of the various possible configurations to be considered in this study, the most grounded (and the most likely to be implemented) is to use the existing roller-shade control algorithm as an upper bound for the shade position. The existing shade control is primarily for the avoidance of glare discomfort and maintenance of views, and may be summarized as follows: (1) keep the shades open as much as possible while still satisfying (2) and (3); (2) do not allow direct solar exposure beyond 3 ft from the facade; (3) and lower shades insofar as necessary to avoid glare.

In this report on initial studies, two different shading configurations are considered: interior venetian blinds and exterior venetian blinds, in both cases controlled continuously for angles between 90^o (horizontal, open) and 10^o (closed). These configurations are the simplest to model in EnergyPlus, and provide insights into the nature of the problem and what to expect (and where to focus further modeling efforts and optimization precisions) with the other configurations.

Because of the inherent complexities of the problem, it is difficult to determine control rules that would minimize the energy use under all possible conditions. So a model-based approach is being used.

3. Methods

3.1 Background on model-based controls for buildings

3.1.1 Online model-based control

Online MPC (Model Predictive Control) offers a way of approaching such problems. Instead of trying to define the control logic explicitly, a building model and an optimization algorithm are used within the control system in real-time to calculate the best setpoint values given the current and predicted conditions. In its general configuration, at each controller time step an optimal sequence of control values over a prediction horizon is calculated, only the first of which is implemented, and at the next controller time step the horizon shifts forward one step and the process is repeated. (Note that a number of variants on this are also possible, such as implementing two or more time steps of inputs and performing the optimization less often, or having different control and prediction horizon lengths.) In cases where prediction is not necessary, the setup is the same but without a prediction horizon.

MPC is widely used in other fields – Qin and Badgwell (2003) note its use in more than 4,000 industrial applications. It was a proven practical technique before it was studied theoretically. Good overviews of the field are available in Morari and Lee (1999) and Mayne et al. (2000).

There have also been a growing number of MPC studies for building systems over the past decade, stemming not from controls research but from building energy simulation research. (See for example Mahdavi (2001), Clarke et al. (2002), Henze et al. (2005), and see the review in Coffey et al. (2010). Some more recent work has come from controls researchers from other fields turning their attention to buildings (e.g. Ma et al. (2010) and Oldewurtel et al. (2010)). Potential for energy savings, demand reduction and performance improvement has been shown with a wide variety of systems, including chilled water storage, radiant slab pre-cooling and integrated HVAC and facade control. And as buildings become more complex the benefits of MPC are expected to become more pronounced.

But MPC is currently far from common practice in building design and operation. It is difficult to use most building simulation tools for this because of their slow run-times and the fact that many do not allow the user to explicitly specify initial state values, and the software used by most controls researchers is unfamiliar to most buildings researchers and practitioners. In addition, online optimization is difficult to implement within existing building control systems, and the fact that the control logic is implicit rather than explicit makes it difficult for system designers to integrate it into their design processes.

3.1.2 Model-based control with offline optimization

For some types of MPC problems, multiparametric programming can be used to solve the problem explicitly, providing a set of control laws that fully cover the conditions space and that exactly replicate control behaviour of online MPC (Bemporad et al., 2002). However, this can only be used with certain types of MPC problems (e.g. linear or switched-linear models with linear or switched-linear objective functions), into which forms this case study would be very difficult or impossible to squeeze. And with the possible exception of Modelica (Wetter, 2009b), this approach would not be possible with any of the commonly used building simulation tools. But the idea of explicit MPC is very appealing for buildings applications, because it would be easier to implement in existing building control systems than online MPC, it would allow for faster annual simulations of the controller in the design phase, and being able to visualize optimal control responses over the full conditions space could provide useful feedback to both the controller design process and to the building and system design process in general.

Methods exist and are being further developed to approximate MPC with offline optimization and using common building simulation tools. Current work by May-Ostendorp and Henze considers the approach of simulating online MPC over some or all of a representative weather year and then using statistical techniques to derive near-optimal control laws from the results. This could provide a useful way of getting these benefits. The approach used herein is similar but slightly different: define a grid of conditions (initial states and predicted disturbances) that covers the range of what the system will face, solve the MPC optimization problem at each point in the grid, and then use the resulting grid of optimal control responses as an interpolation lookup table in real-time control. This approach is described in general in the following subsection, and in greater detail and with a variety of other case studies in (Coffey, 2011).

3.2 Software and methods for offline optimization over a grid of conditions

Open-source software for MPC with standard building simulation tools was developed in previous research (Coffey et al., 2010), using GenOpt (Wetter, 2009a) as the optimizer, which allows for the use of any text-file based building simulation tool that can be called from the command line. Figure 2 shows the same basic structure being used to calculate control lookup tables.

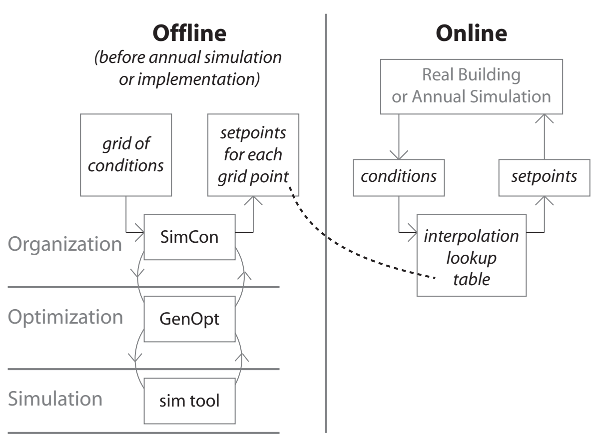

Fig. 2. **Offline optimization over conditions grid**

Given a model of the building and/or system of interest, the steps involved in using this process are:

- Create template input files by demarcating the control variables (with \%) and conditions variables (with \$) in the model and weather files, and locate the objective function in the output file.
- Define bounds and precisions for the control variables, and define conditions variables to be used in the grid (this may be a smaller set than the conditions variables used in the model, as discussed below) along with bounds and spacings for the grid. Configure GenOpt and SimCon with these.
- Use SimCon and GenOpt to solve for the optimal control values at each point in the conditions grid. This is usually a computationally expensive process, but is easily parallelized, and with the use of many virtual machines on now easily-accessible cloud computing platforms, the question is more about money than about time.
- The resulting grid of optimal control values can then be used as an interpolation lookup table for control in simulated or physical implementations, and the multidimensional grid can be visualized through scatter plots or by graphing 2- or 3-dimensional slices through it, providing important feedback to the design process.
- This process can be used iteratively. It is wise to start by solving with a very coarse grid, and possibly with fewer conditions or control variables, fixing model bugs or adjusting objectives and bounds based on the feedback from the grid visualizations, and then building up the precision of the lookup table over 3 or 4 (or possibly more) iterations, keeping grid point solutions from the previous iteration if no changes were made.

3.2.1 Range of applicability

The basis for Figure 3 below is simply that the computation time required for lookup table creation is the product of the conditions grid size and the average computation time per grid point. The dollar costs are based on $0.10 per processor-hour, which is roughly the current commercial cloud computing cost for small-scale users. The shaded area is a conservative cut-off for financial feasibility for a consulting or design firm working on a single building, assuming some iteration in the process. The figure highlights the trade-offs between model complexity, optimization precision and grid spacing, and shows the scale of problems that are feasible - simulation time must be within seconds (note that this is over a simulation horizon of hours or days, rather than a full year), and the number of conditions variables must be kept to less than roughly 5 or 6. As such, this approach usually requires approximations to limit the dimensionality of the lookup table.

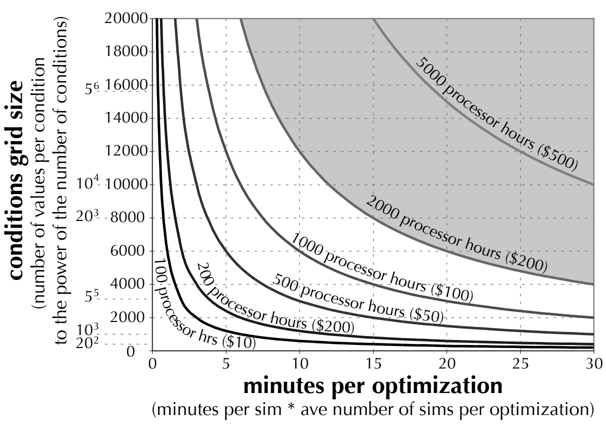

Fig. 3. **Computational requirements**

3.2.2 Approximation techniques

Consider a control problem that requires day-ahead predictions of ambient temperatures, such as the overnight cooling of a massive chilled floor or ceiling. If hourly predictions were used, this would require 24 dimensions in the conditions grid just for the ambient temperature, making the offline solution computationally infeasible. One way of decreasing the dimensionality is to use a coarser prediction, for example using average temperatures for 4-hour blocks instead of 1-hour blocks. Another approach is to take advantage of the expected shape of the curve, and use a small number of parameters to define the prediction, such as the maximum and minimum temperatures and maybe the time of their occurance. In any such variation, a normalized curve is required to produce the values of the temperature at each timestep of the simulation, based on the predicted values of the parameters. Normalized curves can be derived using typical or historical data for the site or system under consideration, as demonstrated below. Similar approaches can be used for other disturbance variables, or to relate disturbance variables to one another.

Even with the approximations of input parametrization, many control problems in buildings still have too many dimensions to be tractable as lookup tables. However, the approach may still be useful in such a case, if the structure of the problem allows it to be decomposed into a hierarchical set of problems where some of the subproblems are small enough to be solved offline. This is discussed in detail and used in case studies in (Coffey, 2011).

3.2.3 Open-source software

The SimCon software described in (Coffey et al., 2010) has been extended to be used for this offline approach. It is written in java, and the source code will be freely available for download (site tdb). The software's current functionality includes: the option of running in either online MPC mode or in lookup table calculation mode, so the same software can still be used for online MPC, including cases with decomposed problems that involve a higher-level problem that must be solved online and lower-level problems that can be solved offline as lookup tables; core methods to set up a sequence of optimization problems for a user-defined conditions grid, solve them using GenOpt, and record the results to a lookup table file; multi-variable interpolation for lookup tables stored in a text file; and an extensible library of algorithms to convert conditions inputs to and from parametrized forms. Peripheral components are also included in the open source code, including sample files for running annual simulations with the controller in the Building Control Virtual Test Bed (BCVTB, (Wetter and Haves, 2008)), java methods for collecting weather predictions from the National Digital Forecast Database, sample java-based interfaces for human-in-the-loop implementations, and a visualization tool for the lookup tables and a weather conditions parametrization tool, both currently in Excel-VBA.

3.3 Overview of steps and report structure

The structure of this report mirrors the required steps in the offline-optimization procedure:

- Develop model (Section 4)
- Develop control optimization structure, including conditions parametrization and grid definition (Section 5)
- Compute control lookup table, graph and analyze the results (Section 6)
- Test controller through annual simulations in BCVTB (Section 7)
- Debug, refine model and optimization structure, repeat

4. Model description

An EnergyPlus version 6.0 model of a single floor was constructed, using some of the constructions and schedules information in an existing EnergyPlus version 1.2 model from the design phase, but with new zoning and new UFAD modeling capabilities. Figure 4 shows the model of the full floor. Figure 5 shows a one-perimeter-zone extraction from the full floor model, which is used in the study herein. Table 1 shows some of the key model parameters and their values in the current model - these values and their associated schedules should be calibrated using actual building data that is currently being collected.

The EnergyPlus model uses the EPMacro language extensively, structured as a set of include files (e.g. different files for geometry, internal loads, HVAC distribution, etc) and one main imf file that references these files and also lists key global parameters, such as the control setpoints. This facilitates model debugging and makes parametric analysis, calibration and optimization easier to configure.

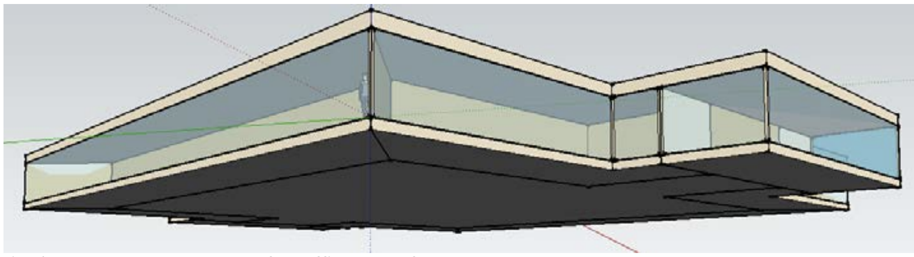

Fig. 4. **EnergyPlus model of an office tower floor**

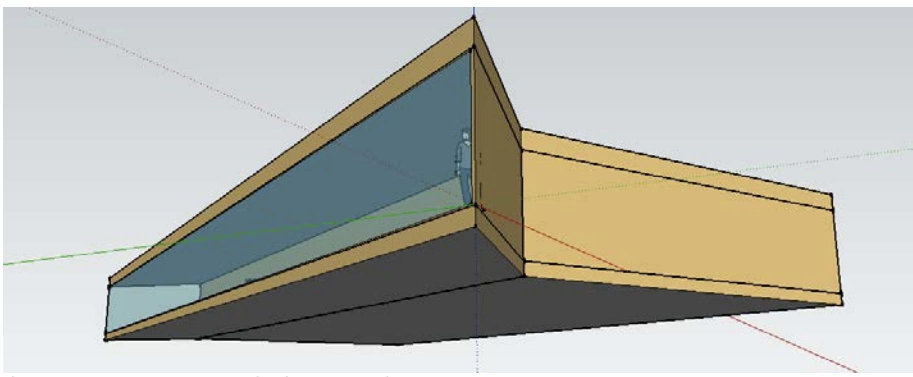

Fig. 5. **Extracted model with just one perimeter zone**

Table 1: Some key model parameter values

lighting power density	4.0 W/m2
equipment power density	10.0 W/m2
occupant density	15 m2 / person
chiller COP at rated conditions	3.0
pumps motor efficiency at rated conditions	0.87
pumps total efficiency at rated conditions	0.62

The control setpoints are the blind angle (between 0^{o}-90^{0}) and the supply air temperature (between 12-17C). In the baseline for comparison, the control values are kept constant at a 90^{o} blind angle and a 12C supply air temperature.

Two different cases were tested: interior blinds and exterior blinds. Figure 6 and Table 2 show the annual energy consumption by end use for the interior blinds case with the baseline control. The end use breakdowns in the exterior blind base case are similar. Note from Figure 6b and 6c that the variations in cooling and fan energy are driven primarily by variations in the ambient temperature, with the solar gains variations playing a secondary role. Also note the decrease in lighting energy use as the direct solar increases from to 200 W/m2, and that the decrease saturates beyond that point.

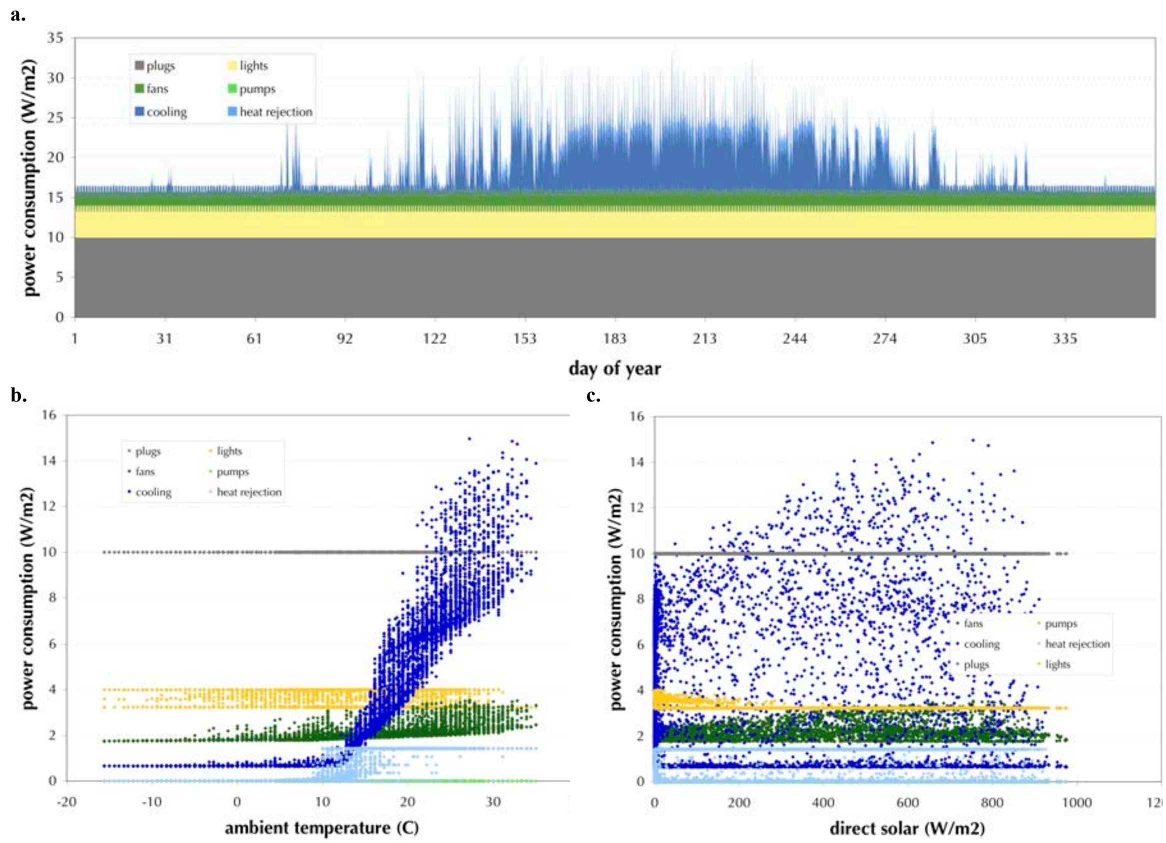

Fig. 6. **Annual energy basecase, end-use breakdowns**

Table 2: Base case annual energy consumption by end use, kWh/m2

	total	lights	plugs	fans	pumps	cooling	heat reject.
base case	170.92	31.79	87.60	17.65	0.03	27.74	6.11
% of total		19%	51%	10%	0%	16%	4%

5. Control optimization configuration

5.1 Assumptions and overview

Because the building envelope is primarily glass and the zone temperatures are kept relatively constant, it is being assumed that thermal mass has little effect on the optimal values of the blind position and supply air temperature. Neglecting thermal mass simplifies the control optimization analysis significantly because it can be done without considering a prediction horizon. The effects of this assumption can be tested later by closer inspection of annual simulations of the resulting controller used in the building with its thermal mass being modeled accurately.

As noted above, the offline optimization approach is generally only feasible with less than roughly 5-6 conditions variables. For this analysis, 5 conditions variables have been chosen a priori for the grid. This selection may be done more rigorously by analyzing the sensitivity of the optimal control values to different conditions variable choices, but that is a laborious process that can hopefully be avoided through good engineering judgment. The 5 conditions chosen in this case are as follows: ambient temperature, direct beam radiation, day of year, diffuse horizontal radiation, and time of day.

The day of year and time of day variables are being used primarily to capture solar position, but may also be used to estimate internal loads in the model with calibrated schedules. Note that the EnergyPlus weather files have many more variables than just the 5 conditions listed above. The remaining variables are estimated as functions of these 5 conditions, as described in the conditions parametrization section below.

The general process for setting up the control optimization and calculating the lookup table is as follows:

1. set up the EnergyPlus model to run over just the desired horizon length and to output just the desired objective function, and demarcate control setpoints and configuration variables appropriately in the EnergyPlus model and the weather file so that they may be used in GenOpt and SimCon
2. define the relationship between the selected conditions variables and the remainder of the conditions needed in the weather file, and code this relationship into the conditions parametrization part of SimCon
3. determine bounds and spacing for the conditions grid, and solve for each point by running SimCon

5.2 Model configuration for control optimization

Although EnergyPlus is generally used for annual simulations, it may also be used for simulation lengths as short as one day. In this analysis, the simulation would ideally only be run for one timestep (usually 15 minutes or less). But since this is not possible, it is run for one day and the objective function output is limited to just the hour of interest by using a schedule in the output definition. The month, date (derived externally by a pre-processing step in SimCon) and hour of day parameters are thus listed in the main imf file and used by the EPMacro language to define the start and stop days for the run period and the schedule used in the output.

The main imf file of the EnergyPlus model also contains the two control variables, which are used by the EPMacro language to set these values in the appropriate places in the model. The GenOpt and SimCon variables are thus demarcated at the top of the main imf file as follows:

```
##set1 SupplyAirTempVal = %supplyAirTemp%
##set1 BlindsAngle = %blindAngle%
##set1 MonthNum = $monthNum$
##set1 DayNum = $dayNum$
##set1 HourNum = $hourOfDay$
```

The EnergyPlus output reports are limited to just one variable:

```
Output:Variable,*,Total Electric Demand,daily,ReportSched;
```

The weather file, on the other hand, requires a heavier hand in its modification. In addition to the monthNum and dayNum values, there are 14 weather variable values that must be entered in the weather file, as shown below. The parametrization to go from the 5 conditions variables to these 14 weather inputs is described in the next section.

```
LOCATION,New York Central Prk Obs Belv,NY,USA,TMY3,725033,40.78,-73.97,-5.0,40.0
...
DATA PERIODS,1,1,Data,Sunday,$monthNum$/$dayNum$,$monthNum$/$dayNum$
...
1987,$monthNum$,$dayNum$,6,0,[long flag value],$Tamb$,$Tdp$,$RH$,101500,$EtHorRad$,$EtDirNorRad$, ...
        $HorIRsky$,$GlobalHor$,$DirectNorm$,$DiffuseHor$,$GlobalHorIll$,$DirectNormIll$,$DiffuseHorIll$, ...
        $ZenithLum$,190,$WindSpeed$,0,0,40.2,77777,9,999999999,110,0.243,0,88,999,999,99
```

5.3 Conditions parametrization

Many of the EnergyPlus weather variables are closely coupled, such as the diffuse horizontal radiation (W/m^2) and the diffuse horizontal illuminance (lux), where one may be reasonably well approximated by a linear correlation with the other. With this in mind, an Excel spreadsheet is used to graph and linearly correlate any particular weather variable with each of the 5 chosen conditions variables, as shown in Figure 7. In this particular screenshot, the zenith illuminance (Cd/m^2) is being compared against (from left to right and top to bottom) the day of year, hour of day, ambient temperature, direct normal radiation and diffuse horizontal radiation. The diffuse horizontal radiation shows a good correlation (shown in closer detail in Figure 8 so the linear curve fit relating the two variables is used in the conditions pre-processor in SimCon to derive the value for zenith illuminance given the diffuse horizontal radiation. For cases where none of the conditions variables correlate well with the particular weather variable (e.g. wind direction, as shown in Figure 9), then the average value for that weather variable is used throughout. Figure 10 shows various graphs from this correlation process, and Table 3 summarizes the correlations that are being used.

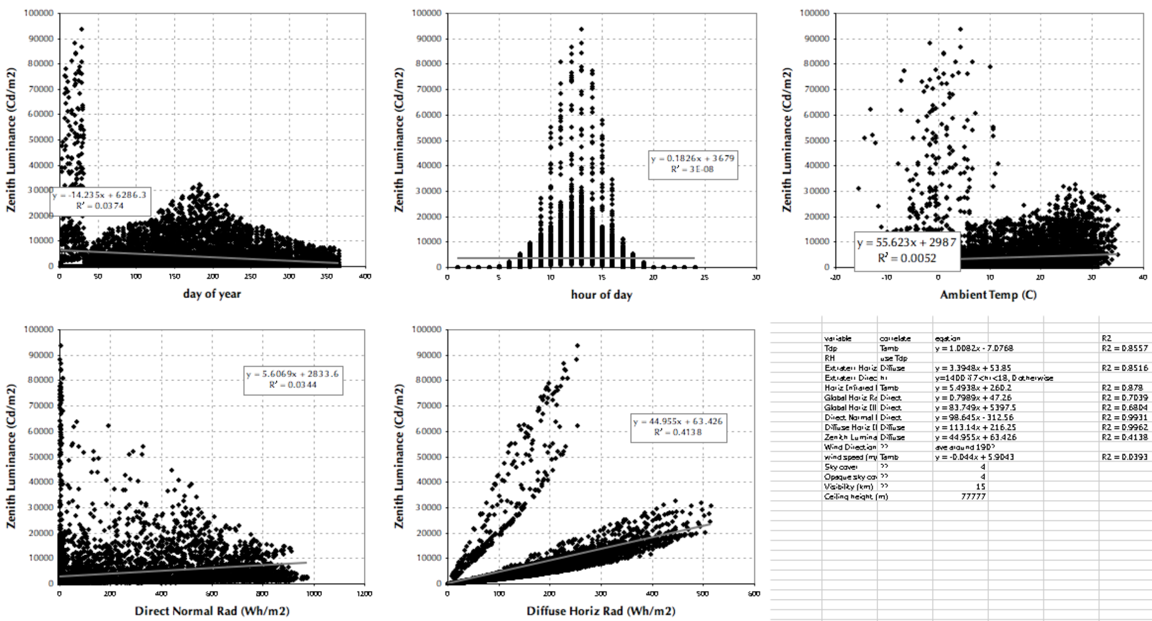

Fig. 7. Screenshot of hourly weather parametrization in Excel, zenith illuminance

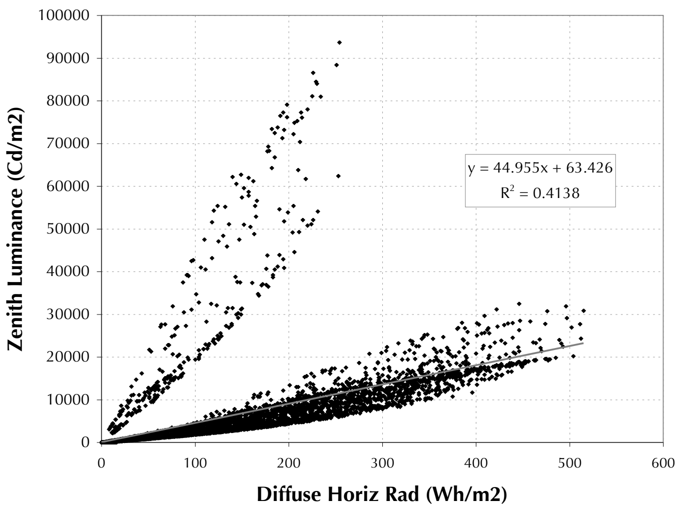

Fig. 8. **Zenith illuminance vs diffuse horizontal radiation**

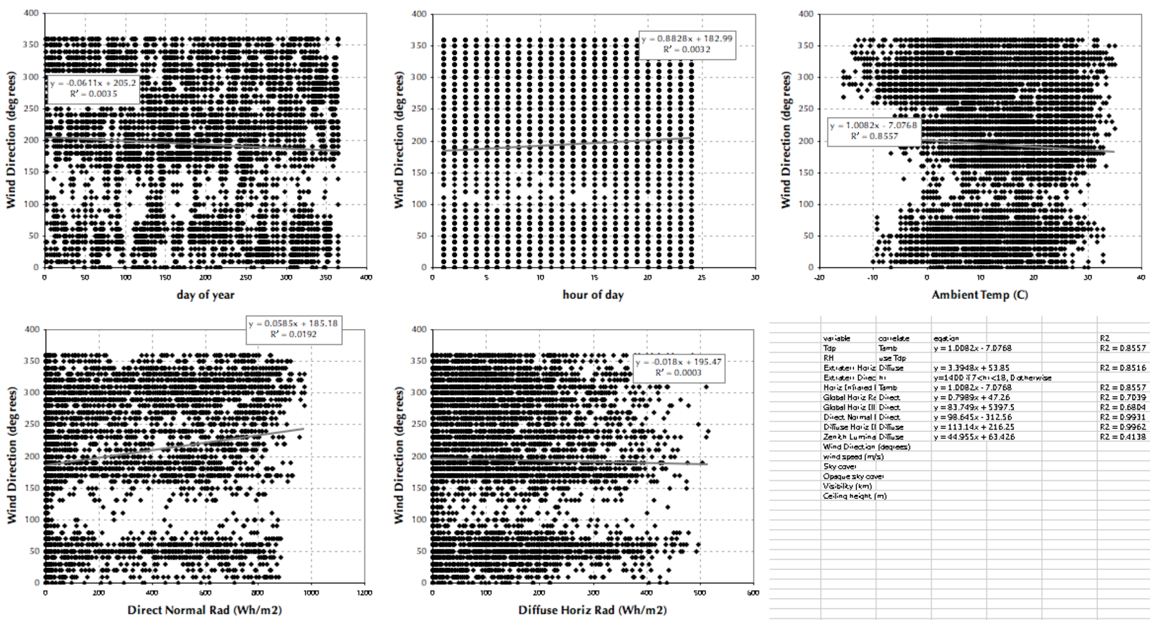

Fig. 9. **Screenshot of hourly weather parametrization, wind direction**

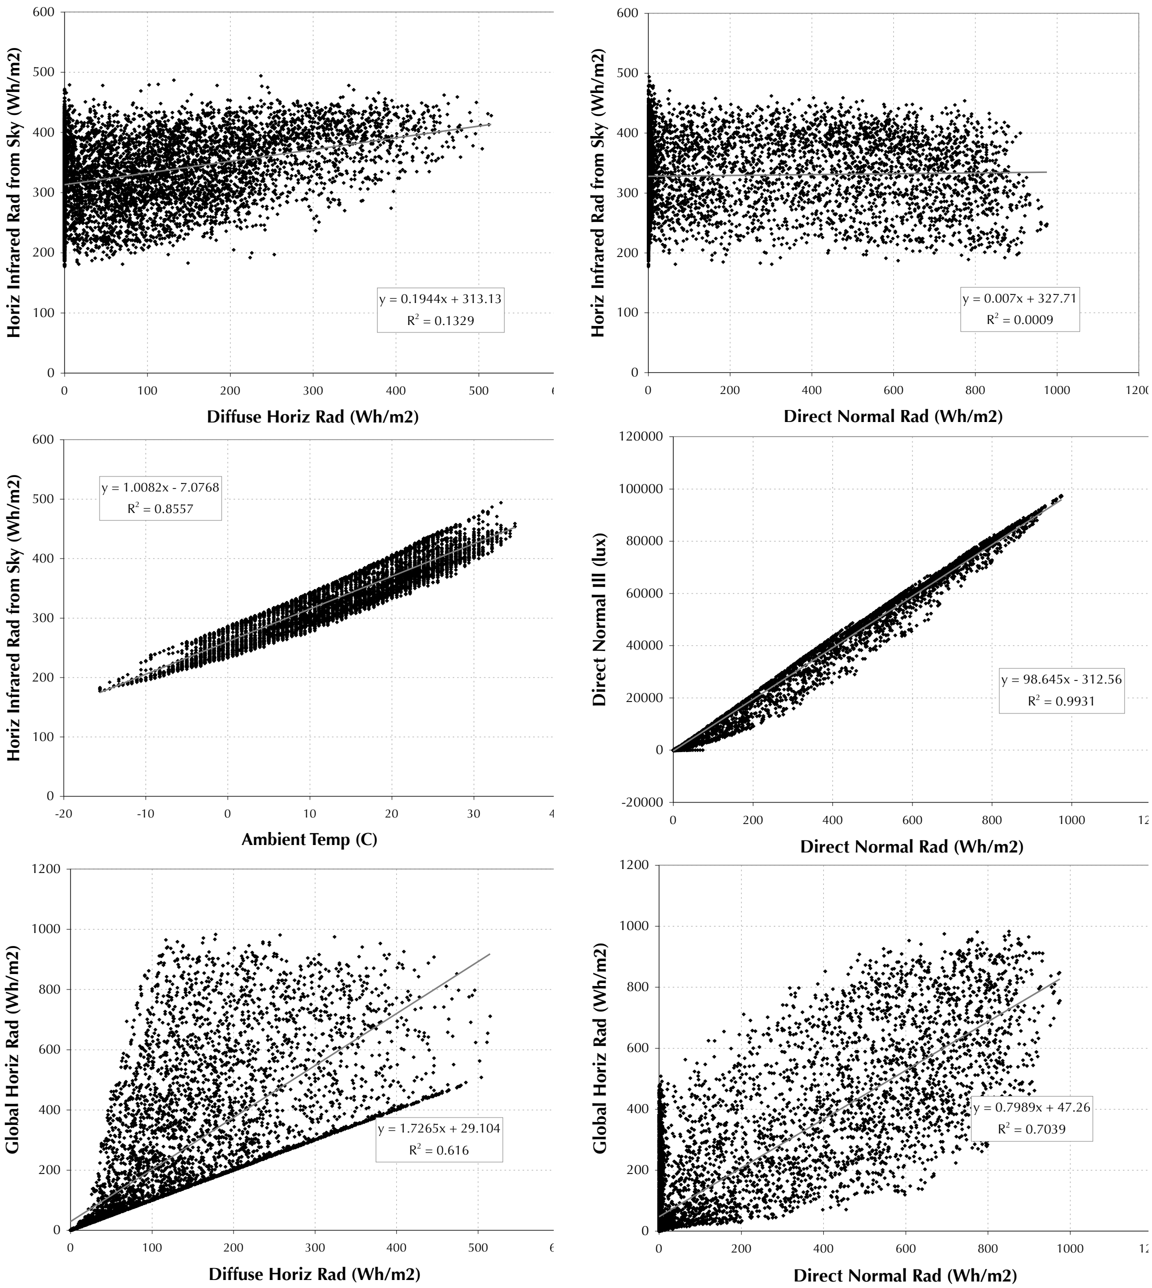

Fig. 10. Example weather variable correlations

Table 3: Parameter correlations

variable	correlate	equation	R2
Tdp	Tamb	y = 1.0082x - 7.0768	0.8557
RH	use Tdp	psychrometric function	
Extraterr Horiz Rad (Wh/m2)	Diffuse	y = 3.3948x + 53.85	0.8516
Extraterr Direct Normal Rad (Wh/m2)	hr	y=1400 if 7$<$hr$<$18, 0 otherwise	
Horiz Infrared Rad from Sky (Wh/m2)	Tamb	y = 5.4938x + 260.2	0.878
Global Horiz Rad (Wh/m2)	Direct	y = 0.7989x + 47.26	0.7039
Global Horiz Ill (lux)	Direct	y = 83.749x + 5397.5	0.6804
Direct Normal Ill (lux)	Direct	y = 98.645x - 312.56	0.9931
Diffuse Horiz Ill (lux)	Diffuse	y = 113.14x + 216.25	0.9962
Zenith Luminance (Cd/m2)	Diffuse	y = 44.955x + 63.426	0.4138
Wind Direction (degrees)	average	190	
wind speed (m/s)	Tamb	y = -0.044x + 5.9043	0.0393
Sky cover	average	4	
Opaque sky cover	average	4	
Visibility (km)	average	15	
Ceiling height (m)	average	77777	

This process can be improved for some variables in future iterations, by considering higher-order and multi-variable correlations. For example, the global horizontal radiation could be directly calculated from the diffuse horizontal and the direct beam by using the solar angle knowing the time of day and day of year.

5.4 Optimization grid configuration

The multi-dimensional grid used for the disturbances is as shown in Table 4. (Note that in the final lookup table used in the controller, the values calculated for day of year = 92 are copied for day of year = 274, and the day of year = 1 values are copied for day of year = 365.) This grid configuration results in 2475 sets of disturbances for which optimal control was to be determined. The Hookes-Jeeves algorithm in GenOpt was used, with 2 step size reductions. The optimization precision was set to 0.25C for the supply air temperature setpoint and 10^{o} for the blind angle setpoint. Approximately 40 processor-hours of computing time were required to solve over the conditions grid (which was carried out twice, once for the interior blind case and once for the exterior blind case).

Table 4: Conditions grid

	min	max	spacing
day of year	1	183	91
hour of day	9	15	3
ambient temperature (C)	5	30	2.5
direct beam radiation (W/m2)	0	400	100
diffuse horizontal radiation (W/m2)	0	400	100

6. Control lookup tables

6.1 Interior blind case

Figures 11 through 13 show some of the many possible slices through the calculated lookup table. Figure 11 shows the optimal supply air temperature as a function of ambient temperature and direct radiation, for various values of diffuse radiation and time of day (all of these graphs are for day of year = 183). The general shape of the optimal supply air temperature curve versus ambient temperature is consistent, with fairly minor deviations for different values of the other variables when the ambient temperature is near 20C. The optimal blind position, graphed in Figure 12 for the same conditions as those in Figure 11, does not show the same consistency. It is particularly noisy when the solar gains are low. But the general trends are as expected, with the blinds tending towards closed when it is hot and sunny, and tending more towards open when it is cold or less sunny. Figure 13 shows the optimal blind position as a function of the direct and diffuse solar gains – here it shows somewhat less noise, and trends towards closed when either the direct or diffuse gains are high, and trends towards open when both are low.

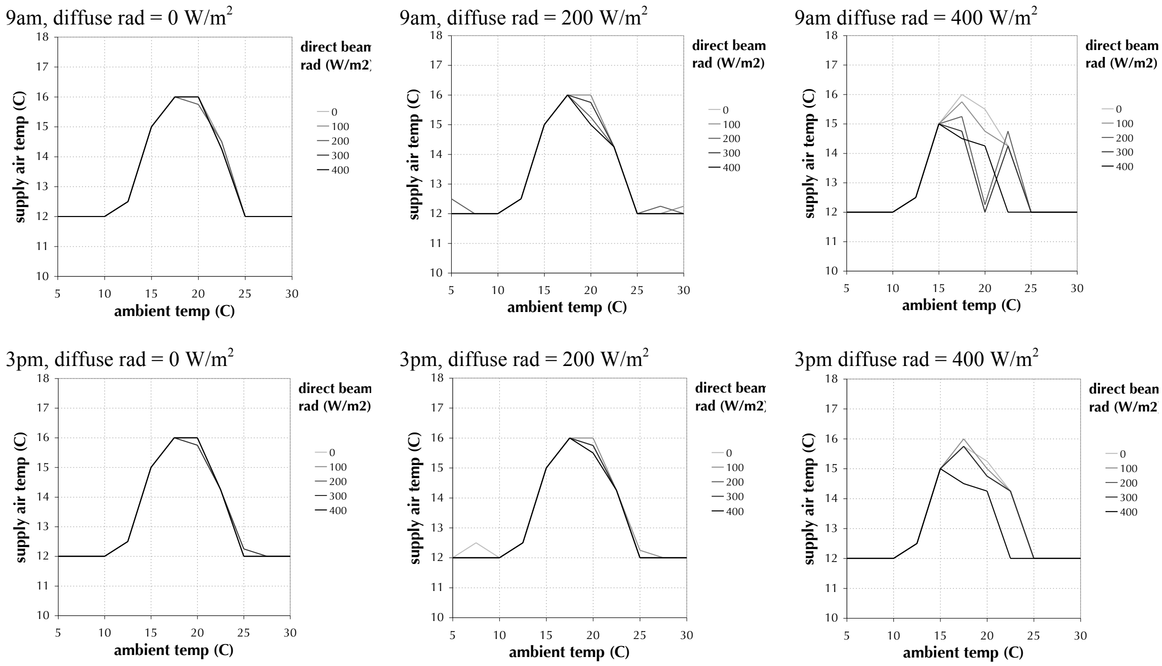

Fig. 11. Interior blind case: Optimal supply temperature under various conditions

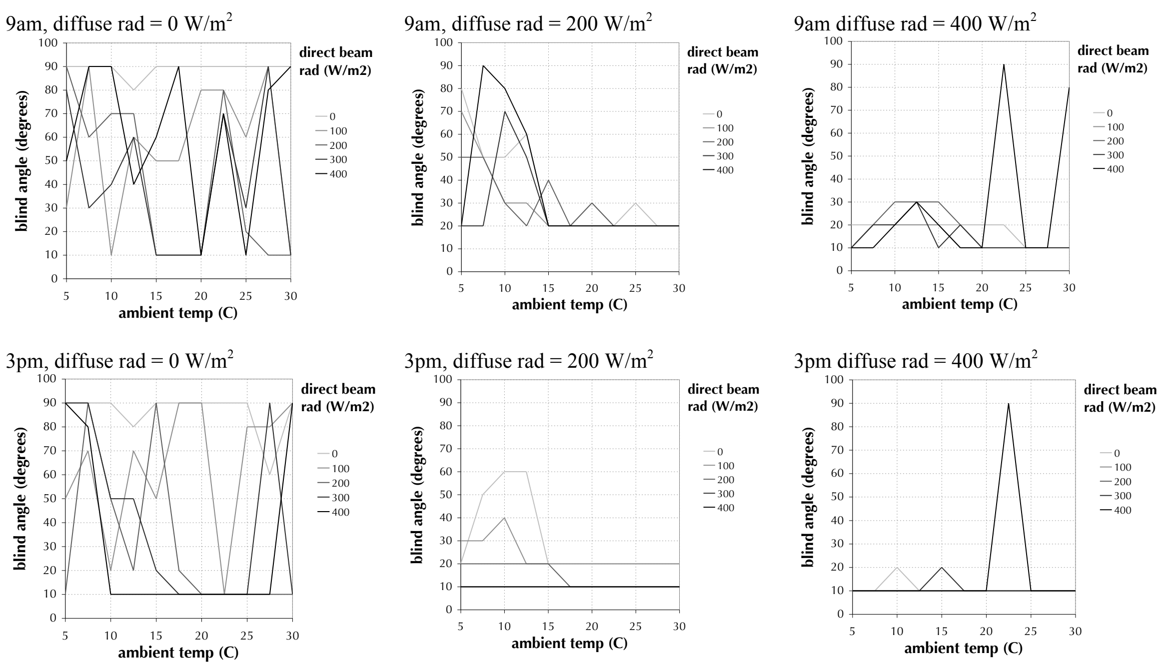

Fig. 12. Interior blind case: Optimal blind angle versus ambient temperature, under various conditions

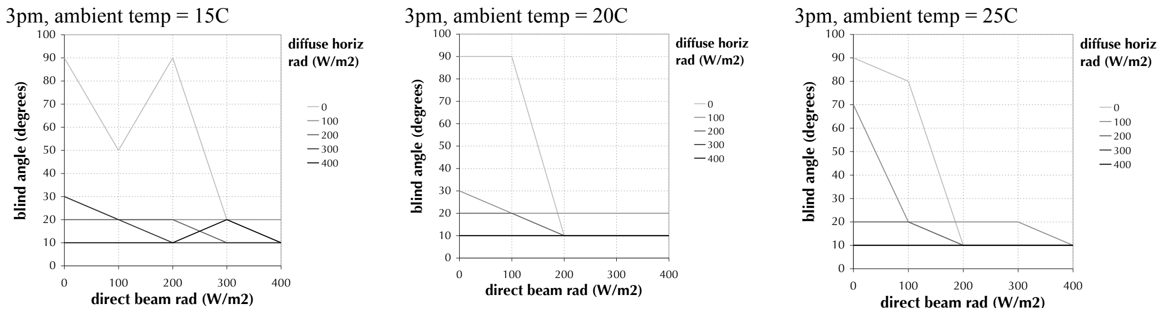

Fig. 13. Interior blind case: Optimal blind angle versus direct and diffuse solar values

6.2 Exterior blind case

Figures 14 through 16 replicate the previous figures but for the exterior blind case. Note that the optimum supply air temperature curve is even more consistent in this case, deviating only very slightly from a constant curve shape versus ambient temperature. The optimal blind position is still noisy at low solar loads, but the trends towards closed under higher solar gains is stronger, and the results also show the optimal blind position to be more responsive to the ambient temperature than it is in the internal blind case.

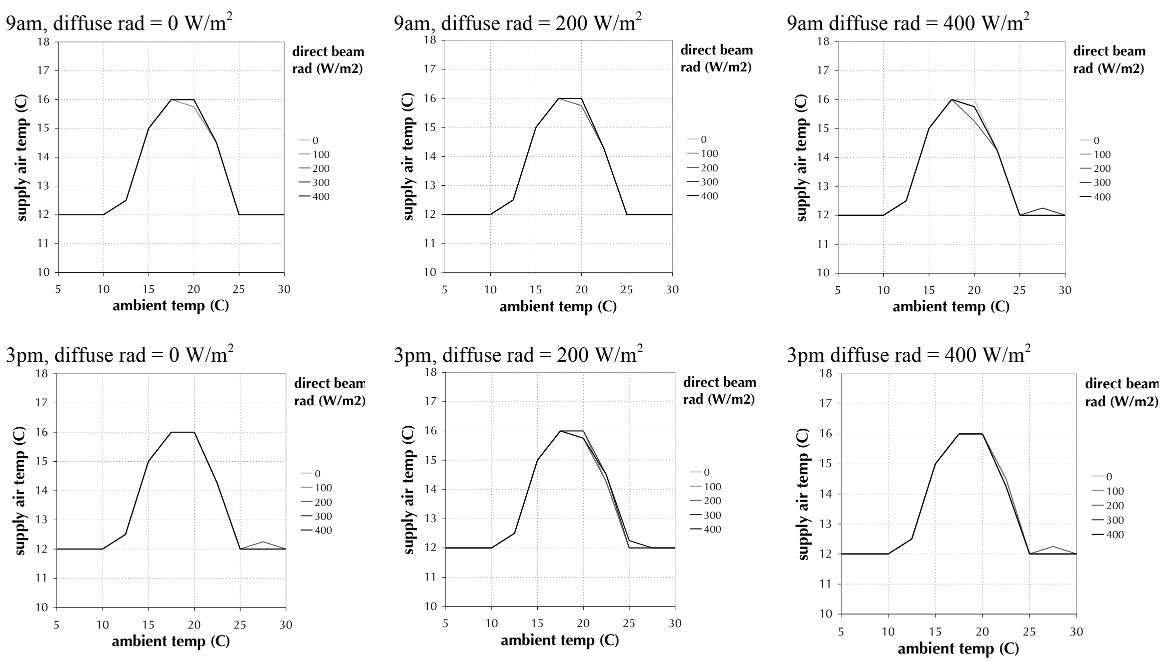

Fig. 14. Exterior blind case: Optimal supply temperature under various conditions

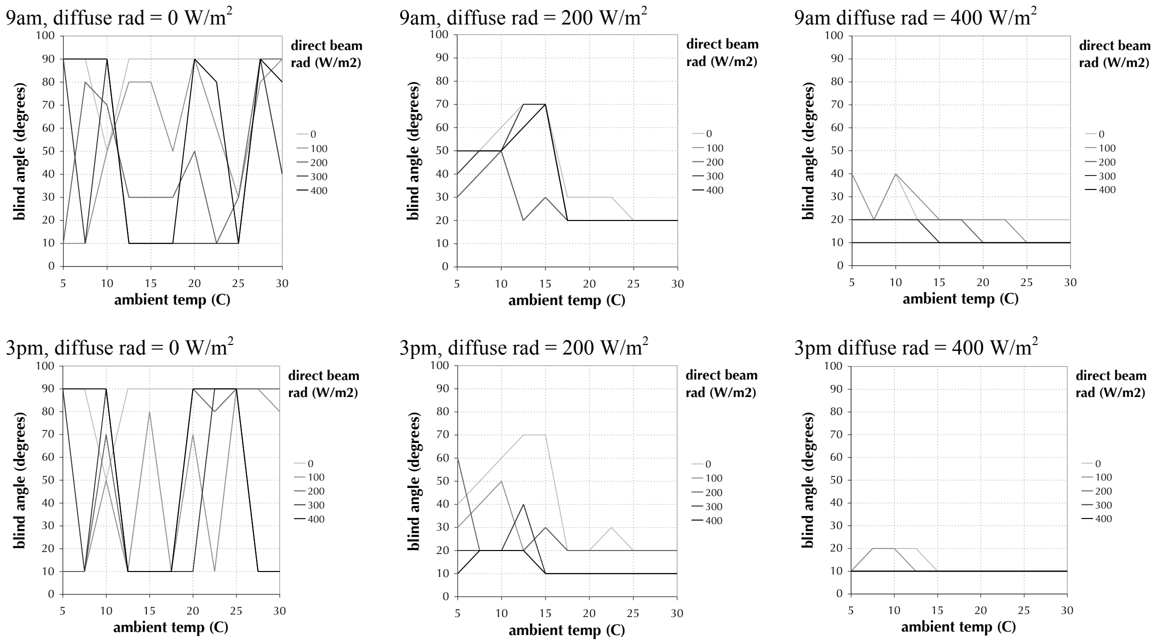

Fig. 15. **Exterior blind: Optimal blind angle versus ambient temperature, under various conditions**

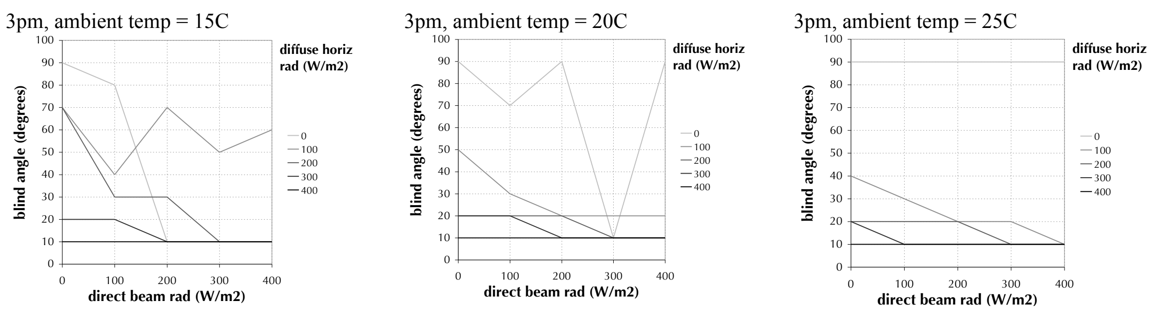

Fig. 16. **Exterior blind: Optimal blind angle versus direct and diffuse solar values**

7. Annual simulations

7.1 Annual simulation configuration

To test the lookup table control, its use is simulated through the Building Control Virtual Test Bed (BCVTB), as shown in Figure 17. The `lookupTableController' module calls a java function that reads the current conditions from a text file and interpolates over the lookup table to determine the control values for that timestep, and outputs those values to text files which are then read back into the BCVTB by the `u1' and `u2' modules. The `simulation' module runs an EnergyPlus model and makes it wait at each timestep for the new setpoints to arrive from the `u1' and `u2' modules. These new sepoints get used in the EnergyPlus model through the Schedule:ExternalInterface syntax in the main imf file. Otherwise, the EnergyPlus model is identical to the one used in the optimizations to derive the lookup table (although in future iterations the effects of model mismatch can be investigated by varying parameters in this annual simulation model). The timestep of the BCVTB simulation is 15 minutes.

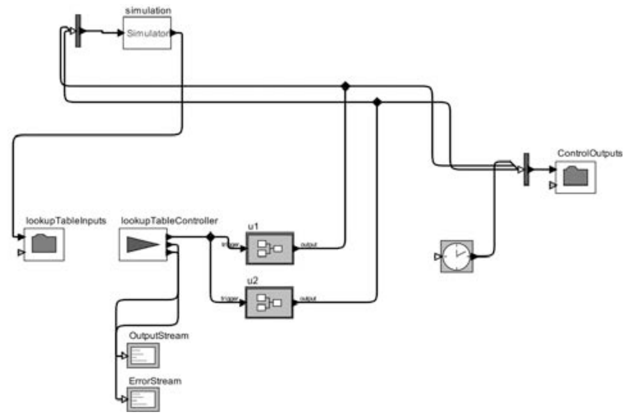

Fig. 17. **Simulated implementation with BCVTB**

The annual simulations were run with the NY Central Park TMY3 weather file. The annual ambient temperatures and solar radiation is shown in Figure 18.

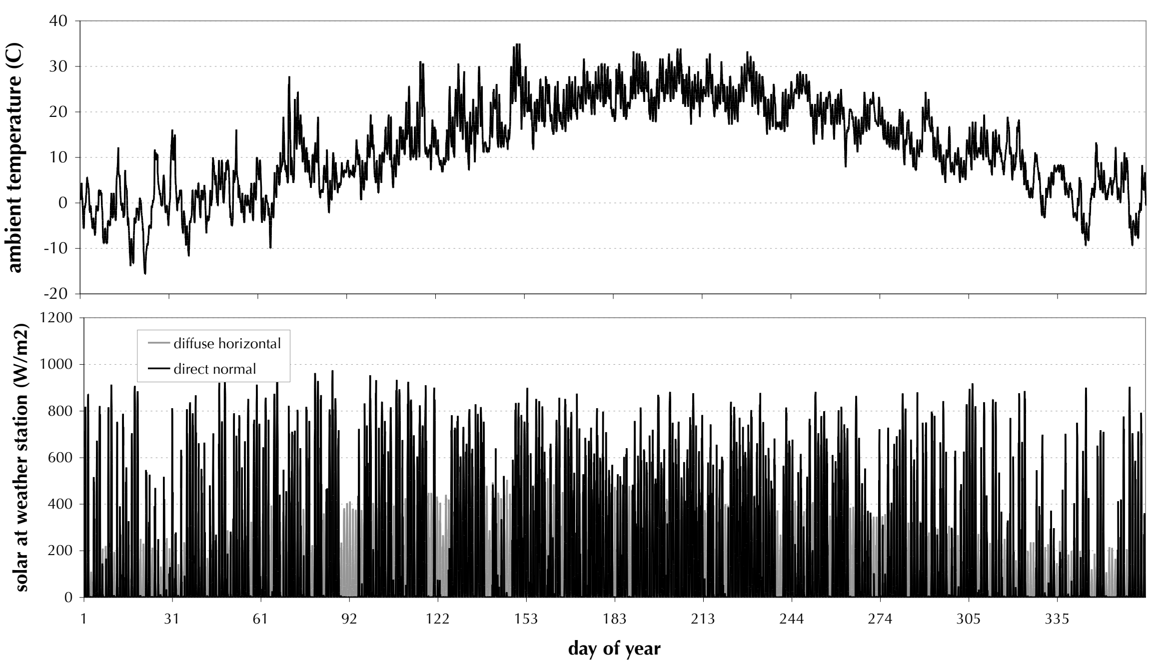

Fig. 18. **Annual weather variables**

7.2 Interior blind case: Annual simulation results

Figure 19 shows the control setpoints over the course of the year, and Figure 20 shows a more detailed view of just the month of March. The x-y plots of the control values versus the ambient temperature (Figure 21) are somewhat more illustrative. The supply air temperature is behaving as expected, and the blind angle is generally trending towards closed when the ambient temperature is higher and open when it is colder, but there is more noise than expected.

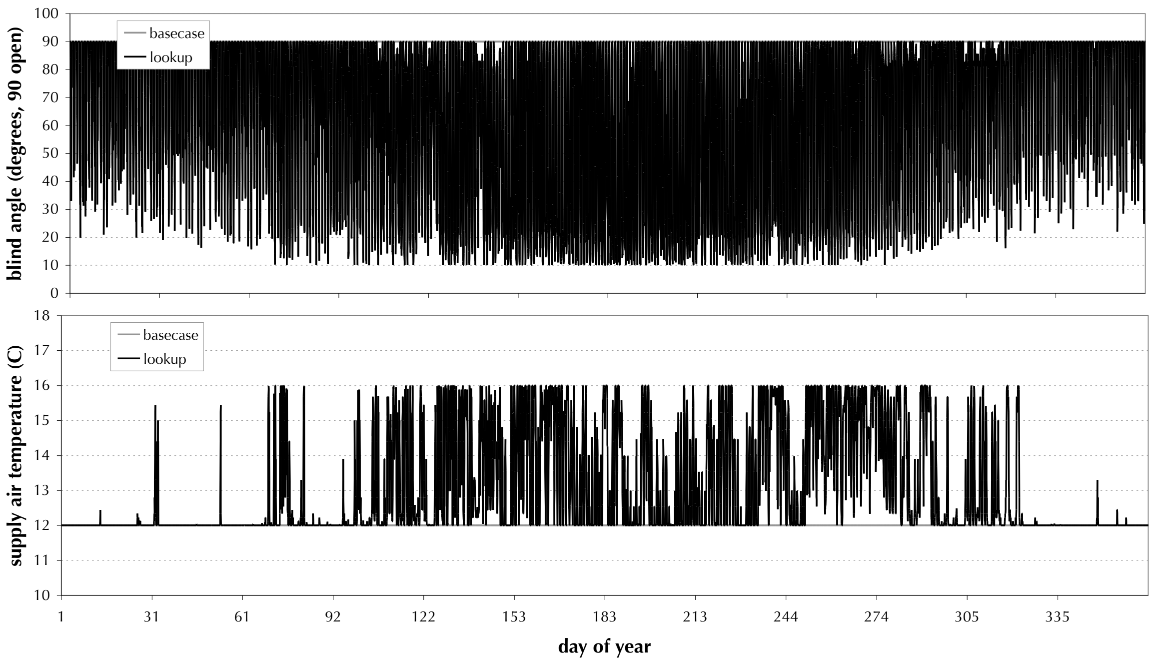

Fig. 19. **Annual control outputs**

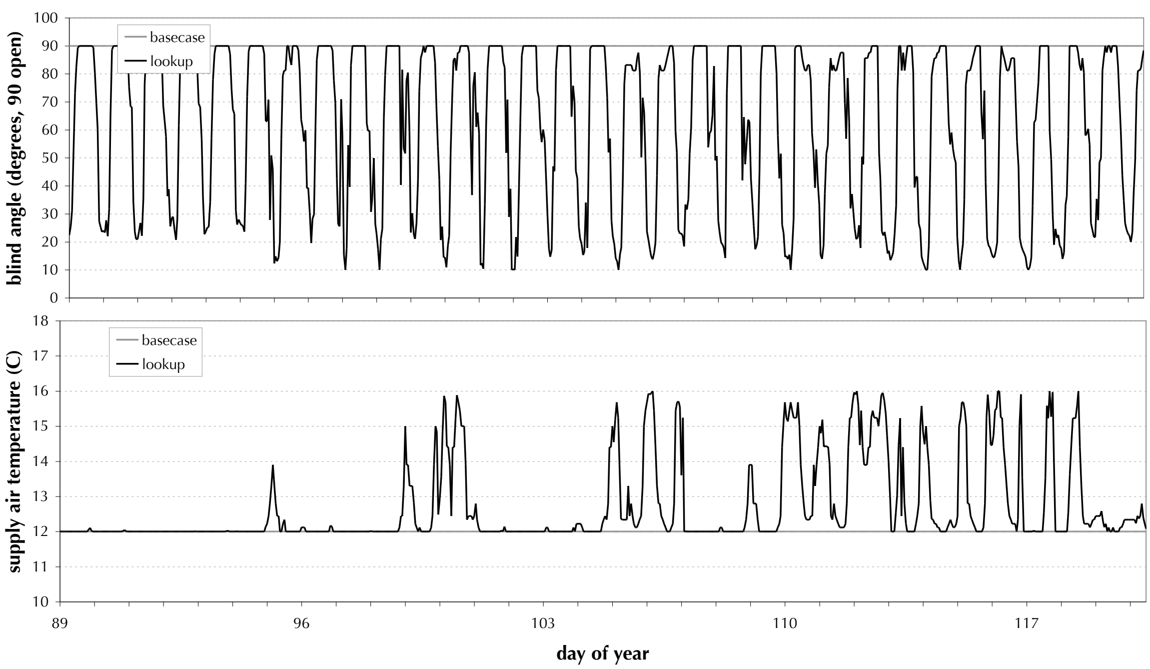

Fig. 20. **March control outputs**

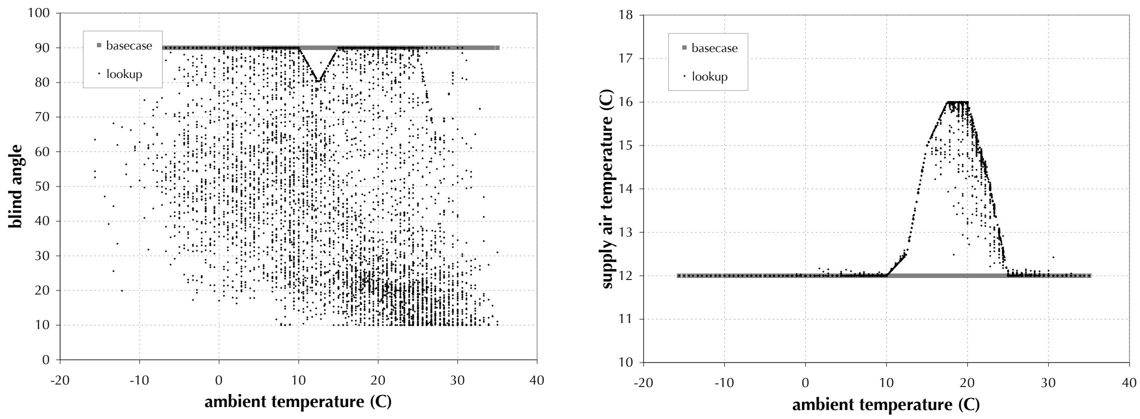

Fig. 21. **Hourly control outputs vs ambient temperature**

Figure 22 shows the energy difference between the lookup table control case and the base case. In general, the lookup table case is using more fan power, but is saving more cooling energy.

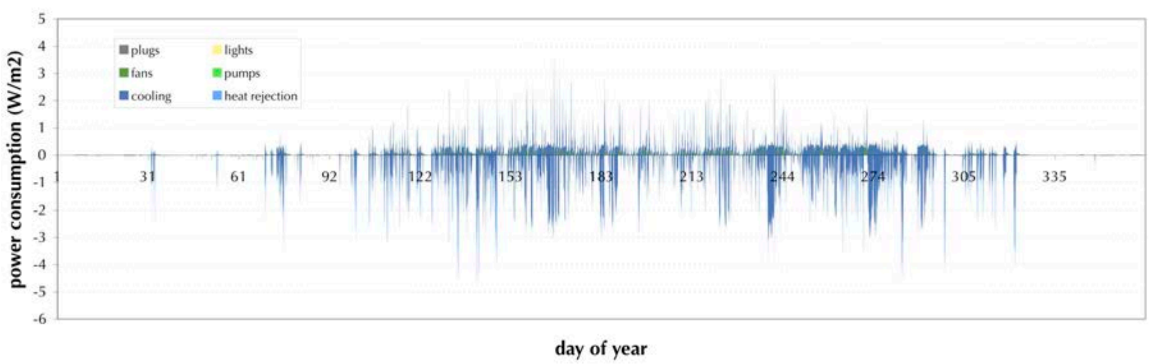

Fig. 22. **Annual energy savings, end-use breakdowns**

The annual energy consumption is summarized in Table 5. The lookup controller is generally decreasing cooling energy compared to the base case by using more shading, and is trading off higher fan energy consumption for lower cooling energy by increasing the supply air temperature. There is a very slight (less than 0.005%) increase in lighting energy in the lookup case. It was suspected that this may be pointing to an error in the daylighting controls in the model, but this was thoroughly checked and is working properly. It seems that the lookup table controller is avoiding increasing lighting use in the way that it is controlling the blind, which would make sense if the cost of increasing lighting use always outweighs whatever cooling savings may be gained by increasing shading.

Table 5: Interior blind case: Annual energy consumption, kWh/m2

	total	HVAC	lights	plugs	fans	pumps	cooling	heat reject.
basecase	170.92	51.53	31.79	87.60	17.65	0.03	27.74	6.11
lookup	168.37	48.98	31.79	87.60	18.54	0.03	24.72	5.69
savings	2.55	2.55	0.00	0.00	-0.88	0.00	3.01	0.42
saving %	1.49%	4.95%	0.00%	0.00%	-4.99%	0.56%	10.87%	6.83%

As with the control values, it is often to view x-y plots of the energy consumption versus the ambient temperature (and against other variables). Figure 23 suggests that the control is not making much difference relative to the baseline when the temperature is less than 10C, and is generally saving energy in the 10-23C range, but seems to be often performing worse than the base case in the 23-30C range, which warrants further investigation, although much of this poor performance is likely caused by the coarseness of the conditions grid and thus will be improved in future iterations (in particular, it is probably caused by the gap between the zero and the lowest solar gains point – this gap has caused similar performance losses in other case studies with this approach, and generally gone away with a finer spacing on solar gains in the conditions grid). Eliminating these areas of poor performance could provide somewhat higher energy savings with the controller.

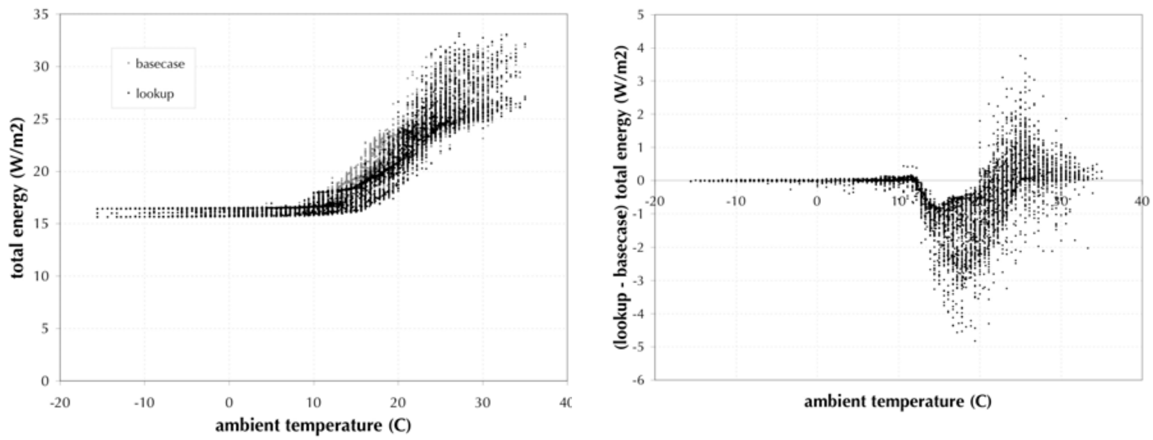

Fig. 23. **Hourly energy consumption vs ambient temperature**

7.3 Exterior blind case: Annual simulation results

The results for the exterior blind case are very similar to the interior blind case, but with slightly higher energy savings compared to the baseline. Table 6 shows the details; the annual HVAC energy savings were 5.56%, rather than the 4.95% in the interior blind case.

Table 6: Exterior blind case: Annual energy consumption, kWh/m2

	total	HVAC	lights	plugs	fans	pumps	cooling	heat reject.
basecase	168.95	49.56	31.79	87.60	17.25	0.03	26.20	6.08
lookup	166.20	46.80	31.79	87.60	18.07	0.03	23.06	5.64
savings	2.75	2.76	0.00	0.00	-0.82	0.00	3.14	0.44
saving %	1.63%	5.56%	0.00%	0.00%	-4.77%	0.58%	11.99%	7.19%

8. Next Steps

8.1 Debugging, further iterations

The results presented here are a second iteration, after lessons learned from a first iteration. Some minor changes should still be made in the model and optimization configuration, and there is still unclear to the author why there is so much noise in the optimal shading position values. Further investigations and iterations are in order. This iterative process of testing and refining is an essential part of the offline optimization approach outlined herein.

8.2 EnergyPlus-Radiance model

The building has two aspects that warrant a more detailed treatment of solar gains than EnergyPlus normally allows. The first is the complex external shading device on the building. The second is the sensitivity of the UFAD system to whether solar gains land on the floor surface or on the furniture. If solar gains land on the floor, that increases the thermal decay in the plenum, producing a higher supply air temperature at the perimeter and a resulting increase in the fan box power consumption.

To deal with these complexities, research modifications were made to the EnergyPlus source code to allow the use of scheduled heat gains on opaque surfaces and window layers, which allows the solar calculations to be carried out externally of EnergyPlus and fed into the thermal calculations. In this case the solar calculations are carried out with Radiance and Window6. After completion of the initial studies using just EnergyPlus, a combined EnergyPlus-Radiance model will be used to provide for more accurate results.

8.3 Testing for different climates and design variants

Once we are comfortable with the results of this study, it can be repeated for other climates and with other design parameter values (such as the reflectivity of the roller shades or the use of active external shading or changes to the building geometries or materials).

9. Value and Prospects of Research Direction

This initial study has shown 5.0% (interior-blinds case) and 5.6% (exterior-blinds case) annual-simulation HVAC energy savings for lookup table controllers for integrated control of shading and UFAD for this particular building and climate. These numbers could increase slightly with the use of a finer grid (particularly in the solar gains). It is expected that the savings for the original case outlined above, with roller shades and glare constraints (the case that would be most likely to be implemented) will produce somewhat less energy savings than those found in these initial case studies. Other configurations may produce slightly more.

The value of the research lies in the following:

- the approach is general enough that it can be used with many different types of systems, and provides a way of dealing with ever more complex integrated systems (such as the system considered here with the addition of natural ventilation or massive slab radiant cooling)
- the approach is such that it can be streamlined and made available to building designers and consultants so that energy savings through better integrated control (or more for more complex systems) can be made commonplace
- the potential savings are greater with more complex system integration, which seems to be the direction that the industry is heading
- perhaps most importantly, the techniques mesh well with existing design processes (through the use of standard simulation tools and the ability to view and understand the controller's behavior) and with processes that should be more commonly used during building operation - fault detection & diagnosis and retrofit analysis with calibrated energy models

The prospects for industry update are significant, if the methods and software are worked through rigorously and disseminated effectively. The success and growth of the company Optimum Energy, which developed out of Tom Hartman's work on near-optimum controls for chilled water systems, is a tribute to the demand. An open-source and more generalized approach, like the one used herein, could find broader use within the industry.

The methods and software are not quite ready for broad use, but could currently be used by leading-edge consulting firms on particular projects (two such firms have already expressed interest in using it). The main methodological issues left to cover are the use of irregular grids for the conditions grids (which could be considered in the next iteration of this study), and devising better ways of working around the problem of implicit state initialization in many simulation tools (which is not a problem in this particular case study because of the negligible mass assumption). Otherwise the main challenges for market uptake are software usability and practitioner awareness and education, both of which can be improved through rigorous testing, documentation and dissemination of this case study and others like it.

10. References

A Bemporad, F Borrelli, and M Morari. Model predictive control based on linear programming - the explicit solution. IEEE Transactions on Automatic Control, 47, 2002.

J Clarke, J Crockroft, S Conner, J Hand, N Kelly, R Moore, T O'Brien, and P Strachan. Simulation-assisted control in building energy management systems. Energy and Buildings, 34(933-940), 2002.

B Coffey. Using building simulation and optimization to calculate lookup tables for control. PhD Thesis, University of California Berkeley (due to be filed in Autumn 2011), 2011.

B Coffey, F Haghighat, E Morofsky, and E Kutrowski. A software framework for model predictive control with GenOpt. Energy and Buildings, 42:1084–1092, 2010.

G Henze, D Kalz, S Liu, and C Felsmann. Experimental analysis of model-based pre- dictive optimal control for active and passive thermal storage inventory. HVAC&R Research, 11(2):189–213, 2005.

Y Ma, F Borrelli, B Hencey, B Coffey, S Bengea, P Haves, A Packard, and M Wetter. Model predictive control for the operation of building cooling systems. IEEE Trans. Control Systems Technology, PP(99):1–8, 2010.

A Mahdavi. Simulation-based control of building systems operation. Building and Environment, 36:789–796, 2001.

D Mayne, J Rawlings, C Rao, and P Scokaert. Constrained model predictive control: Stability and optimality. Automatica, 36(6):789–814, 2000.

M Morari and J Lee. Model predictive control: past, present and future. Computers and Chemical Engineering, 23:667–682, 1999.

F Oldewurtel, D Gyalistras, M Gwerder, C Jones, A Parisio, V Stauch, B Lehmann, and M Morari. Increasing energy efficiency in building climate control using
27weather forecasts and model predictive control. Clima - RHEVA World Congress, Antalya, Turkey, 2010.

S Qin and T Badgwell. A survey of industrial model predictive control technology. Control Engineering Practice, 2003.

M Wetter. GenOpt: Generic optimization program. http://gundog.lbl.gov/GO/, 2009a.

M Wetter. Modelica-based modeling and simulation to support research and development in building energy and controls systems. Journal of Building Performance Simulation, 2(2):143–161, 2009b.

M Wetter and P Haves. A modular building controls virtual test bed for the integration of heterogeneous systems. SimBuild 2008, pages 69–76, 2008.